A Bayesian model for microarray datasets merging

Marie-Christine Roubaud and Bruno Torrésani*

Aix-Marseille Université, CNRS, Centrale Marseille, LATP, UMR 7353, 13453 Marseille, France

April 2012

arXiv:1510.07850v1 [stat.ME] 27 Oct 2015

The aggregation of microarray datasets originating from different studies is still a difficult open problem. Currently, best results are generally obtained by the so-called meta-analysis approach, which aggregates results from individual datasets, instead of analyzing aggregated datasets. In order to tackle such aggregation problems, it is necessary to correct for interstudy variability prior to aggregation. The goal of this paper is to present a new approach for microarray datasets merging, based upon explicit modeling of interstudy variability and gene variability.

We develop and demonstrate a new algorithm for microarray datasets merging. The underlying model assumes normally distributed *intrinsic gene expressions*, distorted by a study-dependent nonlinear transformation, and study dependent (normally distributed) observation noise. The algorithm addresses both parameter estimation (the parameters being gene expression means and variances, observation noise variances and the nonlinear transformations) and data adjustment, and yields as a result adjusted datasets suitable for aggregation.

The method is validated on two case studies. The first one concerns *E. Coli* expression data, artificially distorted by given nonlinear transformations and additive observation noise. The proposed method is able to correct for the distortion, and yields adjusted datasets from which the relevant biological effects can be recovered, as shown by a standard differential analysis. The second case study concerns the aggregation of two real prostate cancer datasets. After adjustment using the proposed algorithm, a differential analysis performed on adjusted datasets yields a larger number of differentially expressed genes (between control and tumor data).

The proposed method has been implemented using the statistical software `R` [1], and BIOCONDUCTOR packages [2]. The source code (valid for merging two datasets), as well as the datasets used for the validation, and some complementary results, are made available on the web site

`www.latp.univ-mrs.fr/~mcroubau/MicroarrayMerging`

1 Introduction

Microarray technology provides high throughput measurements of messenger RNA levels of thousands of genes in tissue samples. Since its introduction almost twenty years ago, it has found applications in many aspects of molecular genetics and functional genomics, from the discovery of basic biological mechanisms to the classification of diseases into subgroups and the prediction of disease outcome.

[1] `www.r-project.org`, the R Project for Statistical Computing.
[2] `www.bioconductor.org`, Bioconductor, opensource software for bioinformatics.

However, although important progress has been made, that has turned lab experiments into industrially standardized protocols, the technology is still facing significant reproducibility problems (see e.g. the introduction of [5] for a more detailed account). Therefore, comparing results from different studies is generally difficult, and extreme care is needed if one wishes to merge different datasets into a single one to improve reliability and generalizability of the results.

Several microarray study aggregation techniques have been proposed. Among them, two approaches have received significant attention. The first one (see for example [4]) proposes to replace numerical expression data with ranked data, in order to correct for study-dependent normalization effects. This implicitely assumes that the ranks are essentially respected in the considered datasets; unfortunately, such an assumption fails to be true in many situations, for example when a large number of weakly expressed genes are present, the fluctuations of which induce large deviations in ranks. The second classical approach, termed *Meta-Analysis*, aggregates the results of statistical analyzes from the different datasets rather than aggregating the datasets themselves (see for example [1] for a review). It has been argued by various authors (see for example [1, 2]) that such meta analysis approaches generally outperform classical statistical analysis directly performed on aggregated datasets.

We propose in this paper a new approach for pre-processing gene expression datasets prior to aggregation, as an alternative to meta-analysis based techniques. The idea is to start with an explicit model for gene expression data, and an explicit model for the distortions induced by the different experiments to be merged. In a Bayesian framework, those two models are used as priors. Assuming normally distributed additive observation noises leads to a closed form expression for the posterior probability.

This model is simple enough to yield a fairly standard estimation algorithm, exploiting elementary linearization and optimization techniques. The latter yields estimates for the model parameters, together with adjusted gene expression values in which both the non-linear distortions and the variances inhomogenities have been corrected, therefore suitable for aggregation and further statistical analysis.

The model and the algorithm are validated in both synthetic datasets (i.e. real datasets with an artificial distortion and additional noise), for which the method is shown to be able to correct for the distortion and noise (to some extent). We then apply the approach to real prostate cancer expression datasets, for which two original datasets are aggregated after correction.

Preliminary results of this work have been given in conference proceedings [8].

2 Approach

The approach we develop here is based upon an explicit modeling of the gene expressions, and of the distortion introduced by the several studies. For the sake of simplicity we limit ourselves to a simple model for logarithms of gene expressions, assuming independent normally distributed values (with unknown means and variances). We also assume that each study gives rise to a non-linear distortion (called an observation function) of gene expressions, and an additive zero-mean Gaussian white noise, of unknown variance. Our approach mainly relies on two steps. First, estimate the parameters of the model (including the nonlinear observation functions), then estimate adjusted gene expression values, that can be aggregated into a single dataset.

More precisely, we rely on the following model for logarithms of measured gene expression data. Assume we are given several datasets $k = 1, \dots K$, involving $g = 1, \dots N_g$ genes (if necessary, after restriction to a suitable subset of common available genes). For each dataset k, several arrays (hereafter termed *conditions* $c = 1, \dots N_c^k$ are available. We denote by y_{gc}^k the logarithms of measured gene expression levels, which we shall call throughout this paper the *observed gene expressions*. The ingredients of the model are the following:

- The *intrinsic gene expression* values $\boldsymbol{x} = \{x_{gc}^k\}$ are assumed to be independent realizations of i.i.d. normally distributed random variables $X_g \sim \mathcal{N}(\mu_g, \sigma_g^2)$ (whose distributions do not depend

on the study k). We write

$$x_{gc}^k = \mu_g + \delta_{gc}^k \ .$$

- The *observed gene expression* values $\boldsymbol{y} = \{y_{gc}^k\}$ are obtained by a study-dependent nonlinear transformation f_k of the intrinsic expression values, with additive Gaussian noise, $Y_g^k = f_k(X_g) + U^k$, with $U^k \sim \mathcal{N}(0, \tau_k^2)$. Componentwise, the observations therefore read

$$y_{gc}^k = f_k(x_{gc}^k) + u_{gc}^k \ .$$

- The *observation functions :* $\boldsymbol{f} = \{f_k,\ k = 1, \dots K\}$ are assumed to be monotonic, smooth functions ($\int |f_k''(x)|^2\, dx < \infty$), and are given a spline-type prior distribution

$$p(f_k) \sim \exp\left\{-\lambda_k \int_{-\infty}^{\infty} |f_k''(x)|^2 dx\right\} \ .$$

Conditionally to the observation functions f_k and the intrinsic gene expressions x_{gc}^k, the observed gene expressions y_{gc}^k are therefore independent and normally distributed, with means $f_k(x_{gc}^k)$ and variances τ_k^2. Applying the Bayes rule, it is easily seen that the log posterior probability distribution is (up to an additive constant) equal to the sum of three components

$$\mathscr{L}(\boldsymbol{x}, \boldsymbol{f}|\boldsymbol{y}) = \mathscr{L}^{(1)} + \mathscr{L}^{(2)} + \mathscr{L}^{(3)} \ , \qquad \text{with}$$

$$\left\{ \begin{array}{lll} \mathscr{L}^{(1)} & = \sum_{k,g} \mathscr{L}_g^{(1);k} & = -\sum_{k,g} \frac{1}{2\tau_k^2} \sum_c [y_{gc}^k - f_k(x_{gc}^k)]^2 \\ \mathscr{L}^{(2)} & = \sum_{k,g} \mathscr{L}_g^{(2);k} & = -\sum_g \frac{1}{2\sigma_g^2} \sum_{k,c} [x_{gc}^k - \mu_g]^2 \\ \mathscr{L}^{(3)} & = \sum_k \mathscr{L}^{(3);k} & = -\sum_k \lambda_k \int_{-\infty}^{\infty} |f_k''(x)|^2\, dx \end{array} \right. \tag{1}$$

The approach we develop here aims at maximizing numerically the above log-posterior probability, estimating the gene average expressions μ_g and variances σ_g^2, the observation noise variances τ_k^2 and the observation functions f_k, and solve the regression problem for estimating intrinsic gene expression values x_{gc}^k.

3 Methods

We describe in this section the estimation method used in our approach, together with the corresponding algorithm. The first step of the algorithm is the estimation of the model parameters: means, variances and observation function. The chosen approach is an iterative one, each parameter is estimated at the time. After the estimation, adjusted gene expression datasets are available, that can be aggregated and analyzed.

3.1 Estimation

3.1.1 Observation functions estimation

Suppose firstly that the *intrinsic* gene expression values x_{gc}^k are fixed. Then the observation function estimation problem reduces to the minimization with respect to $\boldsymbol{f} = \{f_k \in H^2(\mathbb{R}), k = 1, \dots K\}$ of the quantity

$$\Gamma[\boldsymbol{f}] = \sum_k \left[\left(\sum_g \mathscr{L}_g^{(1);k} \right) + \mathscr{L}^{(3);k} \right]$$

and decouples further into the following K independent optimization problems: for $k = 1, \dots K$, solve

$$f_k = \arg\min_f \left\{ \frac{1}{\tau_k^2} \sum_g \left[y_{gc}^k - f(x_{gc}^k) \right]^2 + \lambda_k \int |f''(x)|^2 \, dx \right\} .$$

The latter are actually smoothing spline regression problems (see e.g. [11] for a detailed account), for which efficient algorithms are available and implemented in standard software packages. It is worth noticing that spline functions being piecewise polynomials, once the spline has been estimated, its derivative is readily available.

Remark 1 *These estimations are performed on a set of genes with small variance across samples in each experiment, termed below invariant gene set. This gene set must also be representative of the range of values of the average gene expressions. Practically, the invariant gene set is determined as follows:*

- *Split the range of values of the observed gene expressions into intervals.*
- *Within each interval, select a given percentage of genes with smallest variance.*

3.1.2 Mean and variances estimation

The average gene expressions μ_g are re-estimated at each step of the algorithm as weighted sample averages of the estimated gene expressions: at iteration t,

$$\mu_g(t) = \frac{1}{KN_c} \sum_{k,c} x_{gc}^k(t) \ ,$$

where $N_c = \sum_k N_c^k$ is the total number of arrays.

The variance component estimation is a difficult task here, as many gene variances σ_g^2 are to be estimated. The I-MINQUE (or REML) approach (see e.g. [6]) is a natural choice, which is unfortunately not suitable here as it produces negative estimates for the gene variances. To estimate the latter, we choose sample estimates from the initialization (i.e. pre-adjusted expression data, see Section 3.1.4 below). To estimate the observation noise variances τ_k^2, we resort to the I-MINQUE approach, applied to the *invariant gene set* (defined in REMARK 1).

Remark 2 *Our simulations (see the* E.Coli *example below) show that when the observation noise is large, the corresponding variances τ_k^2 can be underestimated by our procedure, which in turn results in an overestimation of the gene variances. To overcome this problem, a possible solution is to regularize the gene estimates by adding a positive constant λ^2 to the estimated noise variances: modified estimates of the form*

$$\hat{\tau}_k^2 = \tau_k^2 + \lambda^2$$

are considered. In practical situations, since the observation noise is unknown, the regularization parameter λ can be used as a tuning parameter.

3.1.3 Adjusted gene expression values estimation

Suppose now that the observation functions f_k are known, the minimization of the log posterior probability in (1) leads to the optimization of

$$\Gamma'(\boldsymbol{x}) = \sum_{k,g} \left[\mathscr{L}_g^{(1);k} + \mathscr{L}_g^{(2);k} \right] \ ,$$

which splits into a sum of N_g decoupled terms: the estimation of the intrinsic gene expression values x_{gc}^k reduces to minimizing for each $g = 1, \dots N_g$ the quantity

$$\Phi_g(\boldsymbol{x}_g) = \sum_{k,c} \left\{ \frac{1}{2\tau_k^2} \left[y_{gc}^k - f_k(x_{gc}^k) \right]^2 + \frac{1}{2\sigma_g^2} \left[x_{gc}^k - \mu_g \right]^2 \right\} ,$$

where $\boldsymbol{x}_g = \{x_{gc}^k,\, k = 1, \dots K, c = 1, \dots N_c^k\}$.

Due to the non-linearity of the observation functions f_k, no closed-form expression exist for the solution, and we resort to an iterative numerical algorithm. We assume that the mean μ_g and variances σ_g^2 and τ_k^2 are known, as well as the observation functions f_k. At iteration t, assume that the previous estimate $x_{gc}^k(t-1)$ of the gene expression values is available. A linearization of the observation functions f_k in the neighborhood of $x_{gc}^k(t-1)$ yields the first order approximation $x_{gc}^k(t) = x_{gc}^k(t-1) + \epsilon_{gc}^k$, and

$$f_k(x_{gc}^k(t)) \approx f_k(x_{gc}^k(t-1)) + \epsilon_{gc}^k f_k'(x_{gc}^k(t-1)) ,$$

from which we deduce

$$\Phi_g(\boldsymbol{x}_g(t)) \approx \sum_{k,c} \left\{ \frac{1}{2\tau_k^2} \left[\epsilon_{gc}^k f_k'(x_{gc}^k(t-1)) - (y_{gc}^k - f_k(x_{gc}^k(t-1))) \right]^2 + \frac{1}{2\sigma_g^2} \left[\epsilon_{gc}^k - (\mu_g - x_{gc}^k(t-1)) \right]^2 \right\}$$

The update $x_{gc}^k(t)$ can therefore be obtained by optimizing with respect to ϵ_{gc}^k the above expression, which yields

$$a_{gc}^k \epsilon_{gc}^k = \frac{f_k'(x_{gc}^k(t-1))}{\tau_k^2} \left(y_{gc}^k - f_k(x_{gc}^k(t-1)) \right) + \frac{1}{\sigma_g^2} \left(\mu_g - x_{gc}^k(t-1) \right) ,$$

where we have set $a_{gc}^k = [\sigma_g^2 f_k'(x_{gc}^k(t-1))^2 + \tau_k^2]/\sigma_g^2\tau_k^2$.

Set now

$$\alpha_{gc}^k = \frac{1}{a_{gc}^k \sigma_g^2} = \frac{1}{1 + f_k'(x_{gc}^k(t-1))^2 \sigma_g^2/\tau_k^2} ,$$

Then $0 \leq \alpha_{gc}^k \leq 1$, the bounds being attained in the extreme cases (no noise, or constant f_k). This yields the update rule $x_{gc}^k(t) = x_{gc}^k(t-1) + \epsilon_{gc}^k$, i.e.

$$x_{gc}^k(t) = \alpha_{gc}^k \mu_g + (1 - \alpha_{gc}^k) \left[x_{gc}^k(t-1) + \frac{1}{f_k'(x_{gc}^k(t-1))} \left(y_{gc}^k - f_k(x_{gc}^k(t-1)) \right) \right] .$$

i.e. a weighted average of the mean μ_g and the contribution of observations.

Remark 3 *This is similar to empirical Bayes type update rules, the difference being that the weights depend upon the observations, due to the nonlinearity of observation functions.*

The balance between the contributions of the mean and the observations is controlled by the signal to noise ratio *(SNR)* σ_g^2/τ_k^2. *The larger the SNR, the smaller the* α *parameter and the larger the contribution of the observations: adjusted values are closer to the observed ones. Conversely, for small SNR values, adjusted gene expressions are closer to the means* μ_g.

Remark 4 *It may be shown that this update rule is actually equivalent to a variable step gradient descent method:*

$$x_{gc}^k(t) = x_{gc}^k(t-1) + \epsilon_{gc}^k(t-1) \nabla_{x_{gc}^k} \Phi(\boldsymbol{x}(t-1)) ,$$

where the stepsize depends on the gene g, *the condition* c, *the study* k *and the iteration index*

$$\epsilon_{gc}^k(t-1) = \sigma_g^2 \alpha_{gc}^k(t-1) .$$

3.1.4 Initialization.

The estimation procedures described above are used within an iterative algorithm. The latter needs either initial values for the intrinsic gene expressions, or the observation functions, since only data $\boldsymbol{y}$ is available.

We use an estimate for the reciprocal of the observation functions (the so-called *rectification functions*) provided by the approach described in [7]. The estimation of rectification function is formulated in [7] as another smoothing spline problem: optimize, with respect to the mean gene expressions μ_g and the rectification functions φ_k the quantity

$$\sum_{k=1}^{K}\sum_{g=1}^{N_g}\frac{1}{N_c^k}\sum_{c=1}^{N_c^k}\Big[\varphi_k(x_{gc}^k)-\mu_g\Big]^2+\sum_{k=1}^{K}\lambda_k\int|\varphi_k''(x)|^2\,dx\ .$$

The problem is solved by an iterative algorithm, in the same spirit as the approach described here. Besides the fact that the approach of [7] estimates rectification functions instead of observation functions, the main difference with the current approach is that the latter models observation noise, which is not the case of the former. The approach developed in the present work is more complex, but is able to correct for study-dependent observation noise.

3.2 Algorithm

The proposed approach can be summarized as follows:

- **Initialization:** Start from a first estimate $x_{gc}^k(0)$ for the intrinsic expression values, provided by rectification function estimates (see section 3.1.4). Estimate the gene average expressions μ_g and variances σ_g^2 as in section 3.1.2, and the observation functions as in section 3.1.1.
- **Iteration** t**:** estimates $x_{gc}^k(t-1)$ are available.
    - Re-estimate the gene expressions $x_{gc}^k(t)$ as in 3.1.3.
    - Update the mean gene expressions $\mu_g(t)$ as in 3.1.2.

The output of the algorithm consists in estimates $\hat{\boldsymbol{x}}=\{\hat{x}_{gc}^k\}$ for the expression datasets $\boldsymbol{x}=\{x_{gc}^k\}$, to be exploited for further analyses, together with estimates for the means μ_g, variances σ_g^2 and τ_k^2, and the observation functions f_k.

The algorithm was implemented using the `R` statistical environment, from which we used the smoothing spline function `smooth.spline`. Bioinformatics related functions from the `Bioconductor` package were also used.

3.3 Validation: differential analysis

Even though differential analysis is not the only application we have in mind, it provides a convenient setup for validating the proposed approach. In the numerical results we discuss below, differentially expressed genes were searched for in original (distorted) datasets and adjusted datasets, using t-test with FDR correction for multiple testing. We used functionalities from the *Bioconductor* package, namely the `MTP` function from the `multtest` package.

4 Results and discussion

The proposed approach has been tested and validated on several problems, of increasing complexity. We first verified that the algorithm performs well on articifial data, i.e. random data simulated according to the above model. The results are quite satisfactory, but since there's not much to discuss

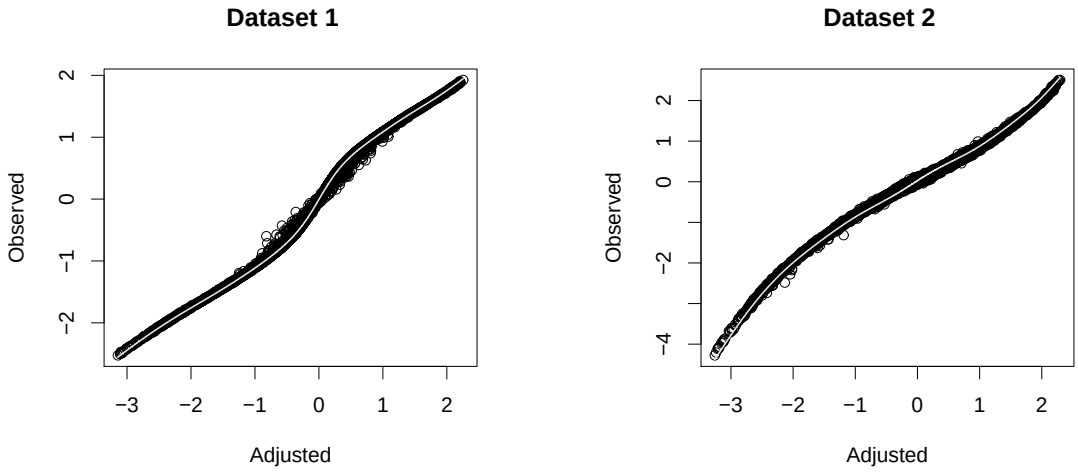

Figure 1: Artificially distorted *E.Coli* dataset: observed data vs adjusted data, and estimated observation functions.

about such simulations we refrain from reporting on them here, and prefer to focus on results closer to practical situations. We report below on two different validations of the approach. The first one concerns *semi-artificial data*, namely real data, with clear biological outcome, which are artificially distorted according to the distortion model above. We show that in such a situation, the proposed approach allows one to almost completely correct for the effect of the distortion, and recover the biological meaning of the data. The second validation is performed on two real human gene expression datasets, focusing on prostate cancer.

4.1 Real data with artificial distortions

A test was performed using artificial observation functions, applied to real data. Namely, we chose a dataset with well understood biological outcome, splitted it into two well balanced subsets denoted by $\mathbf{E_{1o}}$ and $\mathbf{E_{2o}}$ ("o" stands for "original", see below for details) and applied to the two subsets two different observation functions, before adding Gaussian observation noise (yielding distorted subsets $\mathbf{E_{1d}}$ and $\mathbf{E_{2d}}$. After adjustment, datasets denoted by $\mathbf{E_{1a}}$ and $\mathbf{E_{2a}}$ were obtained, that were aggregated. The goal was to study the impact of the deformation induced by the observation functions and the noise (which were chosen so as to hide the biological effects), and the ability of the algorithm to perform a sensible correction.

E. Coli expression data from [3] were used. The data include expressions of 7295 genes under two different situations (20 aerobic and 22 anaerobic). They exhibit a clear variability between the two biological situations, and are thus particularly interesting since they provide a good test of the ability of our approach to recover such a variability after distortion and correction.

Two subsets were created, both involving randomly chosen 10 aerobic and 11 anaerobic conditions. Non-linear transformations f_1 and f_2 were applied to the two so-created subsets (after standardization), and observation noise with variances τ_1^2 and τ_2^2 was added. We report here on the particular case

$$f_1(x) = x^{0.7} \qquad \text{and} \qquad f_2(x) = x^{1.4} \ ,$$

with various choices of the observation noise variances, similar results were obtained with different choices of non-linearities. The values for the noise variances were chosen as a function of the standard deviations of the distorted datasets, following the rules

$$\tau_1 = \alpha \, \mathrm{std}(f_1(\boldsymbol{x}_1))/10 \ , \quad \text{and} \quad \tau_2 = 3\alpha \, \mathrm{std}(f_2(\boldsymbol{x}_2))/10 \ ,$$

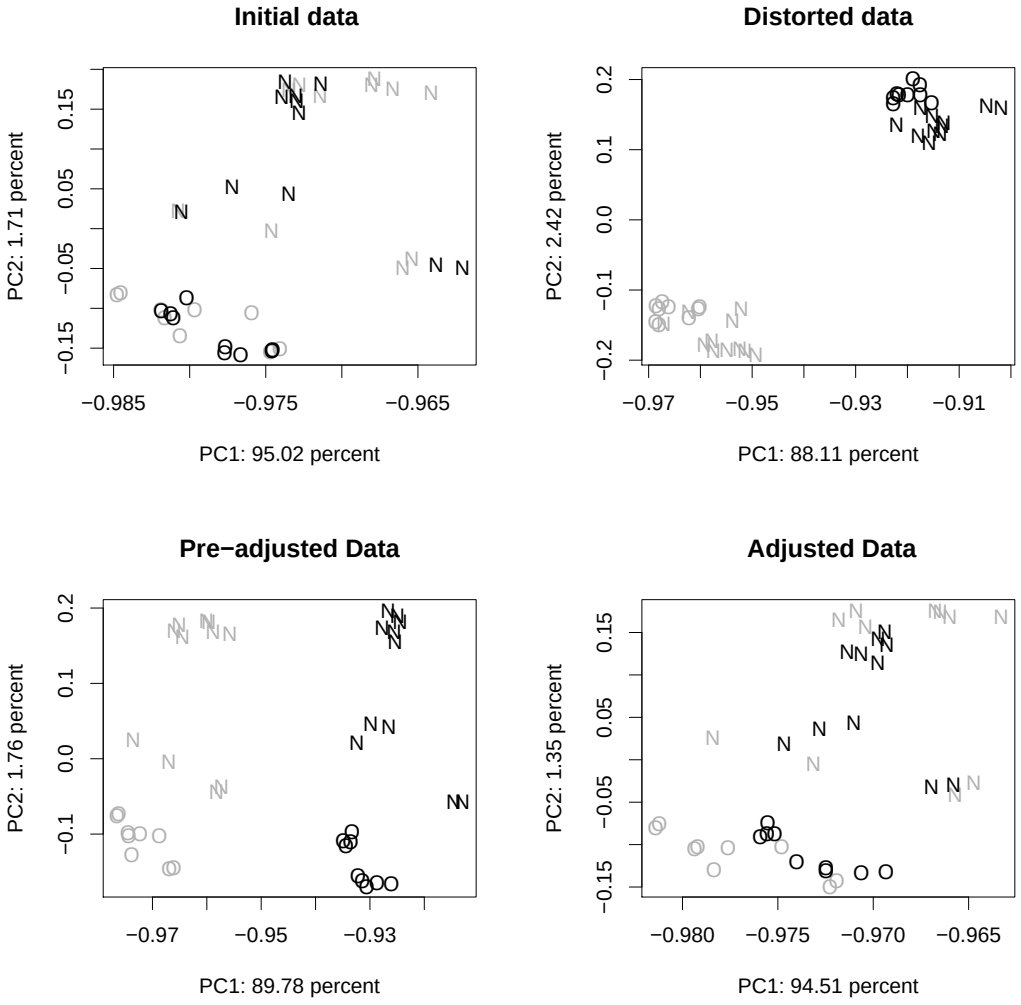

Figure 2: Projections on the first principal plane. Top: original data (left) and distorted data (right). Bottom: rectified data: initialization (left) and processed data (right). O: aerobic; N: anaerobic. Gray: dataset 1; black: dataset 2

	τ_1^2	$\hat{\tau}_1^2$	τ_2^2	$\hat{\tau}_2^2$	$\langle\hat{\sigma}_g^2\rangle$
S1	$2.5\,10^{-3}$	$3\,10^{-3}$	$2.25\,10^{-2}$	$2.1\,10^{-2}$	$1.6\,10^{-2}$
S2	$4.9\,10^{-3}$	$5\,10^{-3}$	$4.41\,10^{-2}$	$4\,10^{-2}$	$1.9\,10^{-2}$
S3	$1\,10^{-2}$	$7\,10^{-3}$	$9\,10^{-2}$	$7.1\,10^{-2}$	$2.6\,10^{-2}$
S4	$1.96\,10^{-2}$	$1.3\,10^{-2}$	$1.76\,10^{-1}$	$1.31\,10^{-1}$	$3.8\,10^{-2}$

Table 1: Observation noise and average gene variance for the four simulated datasets and their estimation.

	O	**D**	**A**	$\mathbf{E_{1o}}\cap\mathbf{E_{2o}}$	$\mathbf{E_{1d}}\cap\mathbf{E_{2d}}$	$\mathbf{E_{1a}}\cap\mathbf{E_{2a}}$
S1+	258	121 (129)	162 (173)	5	1	1
S1-	328	139 (141)	187 (188)	31	9	10
S2+	275	107 (112)	123 (129)	5	2	2
S2-	312	124 (129)	148 (155)	31	3	4
S3+	238	71 (74)	88 (93)	5	1	1
S3-	305	105 (109)	113 (116)	30	3	4
S4+	253	58 (60)	65 (68)	7	1	1
S4-	300	72 (73)	80 (81)	32	1	1

Table 2: Differential analysis on artificially distorted *E.Coli* dataset. The first column indicates the simulation index (see Table 1 for the corresponding observation noise variances), the sign indicates over or under expression in aerobic conditions. Column 2: numbers of differentially expressed genes in the original dataset. Columns 3 to 7: numbers of differentially expressed genes in distorted (**D**) and adjusted (**A**) datasets, and in the intersections $\mathbf{E_{1o}}\cap\mathbf{E_{2o}}$, $\mathbf{E_{1d}}\cap\mathbf{E_{2d}}$ and $\mathbf{E_{1a}}\cap\mathbf{E_{2a}}$ that are present in the original dataset (the numbers between parentheses indicate the total numbers of differential genes).

with various values of the tuning parameter α. Actual variance values, together with the corresponding estimated values, are displayed in **Table 1**, together with the estimated average gene variance, denoted by $\langle\sigma_g^2\rangle$. As can be seen, the estimates are fairly good, except for the last situation, which corresponds to the value $\alpha = 20$, presumably too large for our procedure (results continue to degrade for higher values of α, which probably originates from the degradation of the gene variance estimates).

To illustrate graphically the effect of the adjustment, we display in **Fig. 2** the first factorial plane obtained by (normalized) principal component analysis with original data, distorted data, pre-adjusted data (using the method of [7], used here as initialization, see Section 3.1.4), and adjusted data. The projection of the output data $\hat{x}_{gc}^k$ onto the first factorial plane turns out to reproduce fairly accurately the projection of the original data x_{gc}^k (before distortion). The processing has therefore permitted to recover the biological features as the first source of variability. This is confirmed by visual inspection of the projections on other factorial planes (not shown here).

A more quantitative validation can be obtained through differential analysis. We searched for genes differentially expressed between aerobic and anaerobic conditions, on initial data, distorted data and adjusted data, for various values of observation noise. For the sake of comparison with meta-analysis based approaches, we also performed differential analysis on the two subsets $\mathbf{E_1}$ and $\mathbf{E_2}$ (in original, distorted and adjusted situations) and computed the intersection of the two estimated gene subsets in each situation. The goal of the comparison is to asses the ability of each approach to recover original differentially expressed genes.

The results are displayed in **Table 2**, which gives the numbers of differentially expressed genes on the original dataset (first column), and the number of original differential genes that are recovered on the distorted dataset (second column) and the adjusted dataset (third column); total differential gene numbers are given between parenthesis.

A first conclusion is that the distortion significantly degrades the differential analysis, the degradation being obviously bigger when the observation noise is higher. The processing improves the situation, which is reflected by an average increase of the number of the recovered original differential genes ranging between 11% and 20% (depending on the noise level).

For the sake of comparison, we also display in the last three columns numbers of differential genes obtained by intersecting the results of individual differential analysis on the two sub-datasets, in the three considered situations. It is worth noticing that the loss of differential genes is in this case really huge, which is for us a strong motivation for aggregating datasets rather than intersecting differential analysis results.

4.2 Real data

The proposed approach was also tested on the two publicly available datasets of prostate cancer expression data reported in [9] and [10]. Both datasets originate from Affymetrix HG-U95av2 arrays.

We started from raw expression data, and normalized them separately using the standard RMA normalization procedure from the Bioconductor package. After normalization, the dataset of [9] appeared to exhibit much larger variability than the [10] dataset. We thus performed the following pre-processing: we extracted a subset of arrays whose correlations to the median array exceed 90%, and reduction to common genes). After pre-processing, the two-datasets consist in respectively 61 (32 tumor and 29 normal) and 86 (37 tumor and 49 normal) conditions, with 12625 genes, from which we extracted balanced sub-datasets (29 tumor and 29 normal for each)

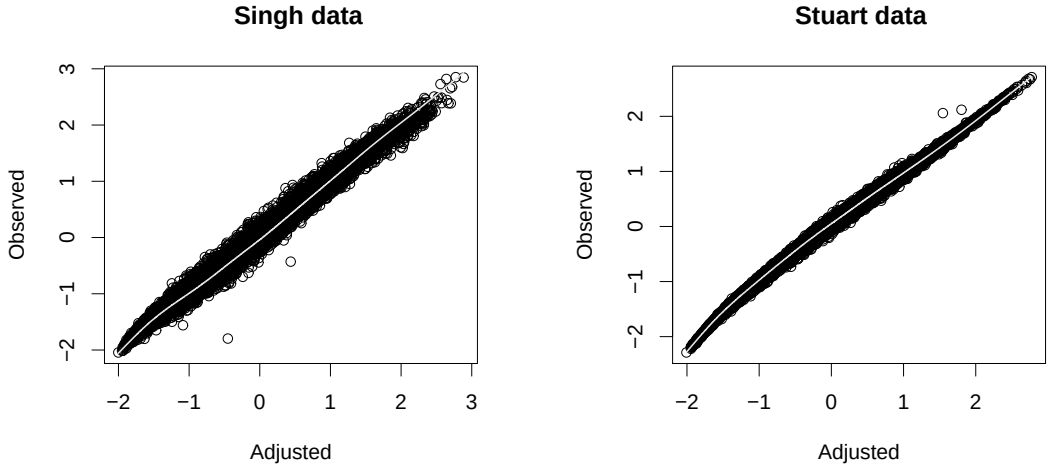

Figure 3: Estimated observation functions for the prostate datasets: Singh (left) and Stuart (right) datasets

The proposed algorithm was run on these two datasets. The result of the processing is shown in Fig. 3 (estimates for the observation functions) and in Fig. 4 (estimated variances, for the observed and adjusted data). We first notice that the estimated observation functions are in this case fairly regular, no large distortion has been found. It can be seen on Fig. 3 that the Singh data (termed $\mathbf{E}_1$) have been shrunk more importantly than the Stuart data ($\mathbf{E}_2$). This can be interpreted in terms of the observation noise. Indeed, the observation noise was found to be significantly larger in the Singh

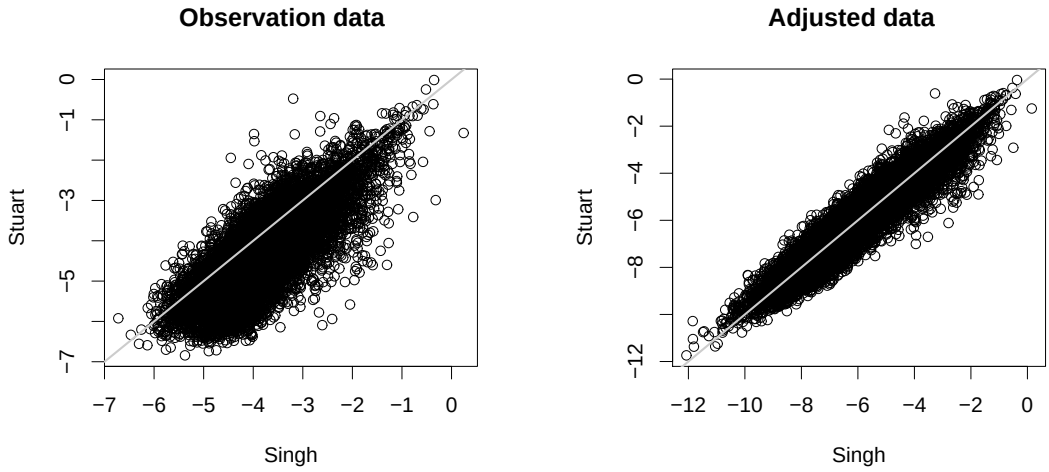

Figure 4: Prostate datasets: estimated gene variances in the original dataset (left) and the adjusted dataset (right), in logarithmic scales

dataset than in the Stuart dataset: $\hat{\tau}_1^2 \approx 2.6\,10^{-2}$ and $\hat{\tau}_2^2 \approx 1.6\,10^{-2}$. Nevertheless, both values are larger than the estimated average gene variance $\langle \hat{\sigma}^2 \rangle \approx 4.2\,10^{-3}$. According to Remark 3, we are here in a situation where the contribution of the average to the adjusted gene expression values tends to be larger than the contribution of the observations.

Finally, it is clear from Fig. 4 that the proposed approach has done quite a good job for adjusting the gene variances: variances in adjusted Singh and Stuart datasets are coherent. The interstudy variability has therefore been considerably reduced.

As in the previous case study, to assess more quantitatively the performances of the proposed approach, differential analysis was performed on the real dataset, and the adjusted dataset. After filtering out the 30% least variable genes, we used t-test based decision rule, with FDR correction for multiple testing ($\alpha = 5\%$, 5000 bootstrap samples). Differentially expressed genes were seeked for the real dataset and the processed dataset, as well as the individual subsets $\mathbf{E_1}$ (Singh) and $\mathbf{E_2}$ (Stuart), and the corresponding adjusted subsets $\mathbf{E_{1a}}$ and $\mathbf{E_{2a}}$. The results are summarized in **Table 3**.

	O	**A**	**O∩A**	**O\ A**	**A\ O**	$\mathbf{E_1 \cap E_2}$	$\mathbf{E_{1a} \cap E_{2a}}$
+	73	88	66	7	22	4	5
-	232	245	217	15	28	1	1

Table 3: Differential analysis on the *Prostate* dataset: numbers of differentially expressed genes. First column: aggregation of the two original datasets. Second column: aggregation of the two adjusted datasets. Columns from 3 to 5: intersection and differences between original and adjusted. Columns 6 and 7: intersections of original subdatasets and adjusted subdatasets respectively.

As before, a striking first result is the fact that the intersection of the differential genes found from $\mathbf{E_1}$ and $\mathbf{E_2}$ individually is extremely small, as well as the the intersection of the differential genes found from $\mathbf{E_{1a}}$ and $\mathbf{E_{2a}}$ individually. Quite a large number of differential genes appear to be lost by this procedure. Besides, as was also seen in the above *E. Coli* simulation study, the adjusted datasets yield a larger number of differential genes than the original datasets. A closer examination of the results of the differential analysis leads to the following conclusions:

- The adjusted dataset features, after concatenation, a larger number of positively and negatively

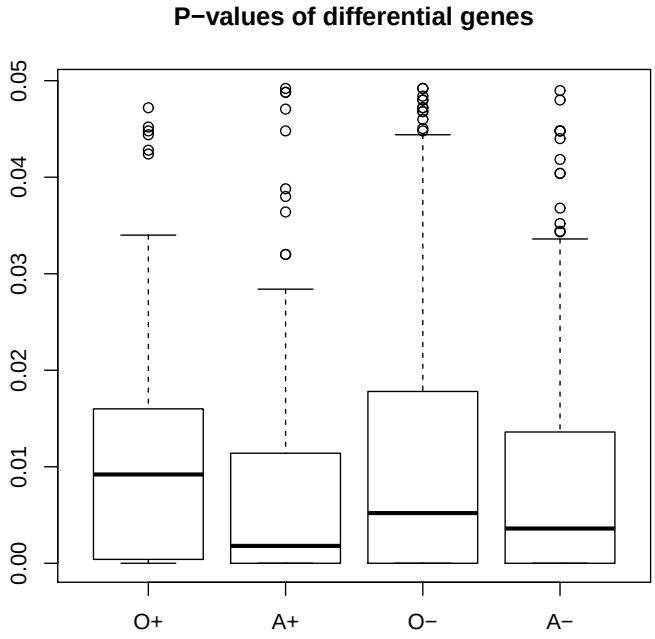

Figure 5: P-values of differential genes for the original and the adjusted datasets: Positively differential genes in the original (O+) and adjusted (A+) datasets, and negativeley differential genes in the original (O-) and adjusted (A-) datasets

differential genes. Even if one focuses on differential genes with small p-values (say, less than 1%), the adjusted dataset yields a significant number of "new" genes, including for example MAPKAPK3, POLR2H, BDH1, PYCR1, PDIA4, SLC25A6, ZMPSTE21, ENTPD6 for the positively differential genes (i.e. overexpressed in tumors), and ADD1, VAT1, PPP2CB, RBPMS, FSCN1, ARHGEF4, GRK5, FEZ1, HTRA1, RBMS1, CETN2 and OSR2 for the negatively differential genes.

A precise analysis of the role of these genes in prostate cancer goes far beyond the scope of the present work. Let us simply notice that most of these genes have already been reported in the prostate cancer literature. This makes them potentially interesting.

- Globally, the p-values of differential genes are smaller in the adjusted dataset than in the original dataset. This appears clearly in **Fig.** 5 where they are displayed in boxplot forms.

- Conversely, less genes were found differential in the original dataset only; in addition, the corresponding p-values are generally close to the critical value chosen in this study (5%). A closer examination shows that their variance in the adjusted dataset is smaller than in the original one. According to Remark 3, this presumably originates from a low signal to noise ratio σ_g^2/τ_k^2, or a large value of the observation function derivative $f'(\mu_g)$ (in the present approach, these quantities govern the parameter α_{gc}^k which controls the variance in the adjusted dataset).

5 Conclusion

We have proposed in this paper a new approach for aggregating gene expression datasets originating from different studies. The rationale of the method is to adjust the datasets prior to merging, using

a Bayesian modeling. The model we develop is specifically adapted to microarray data, and assumes that the inter-study variability originates from non linear distortions, combined with study dependent observation noise. The inter-study variability model is combined with a Gaussian model for the gene expression values (after logarithmic transformation) in a Bayesian framework.

Our results on artificial and real data show that the method is sound, and capable of correcting for study-dependent distortion and observation noise. For the sake of simplicity, we focused on the case of two datasets to be merged, however the model is versatile enough to accomodate arbitrary numbers of datasets. The results on real data discussed here are quite encouraging, however the approach will have to be tested further on a larger scale to be validated more thoroughly.

Several aspects are to be investigated further. Among them, the problem of variance components estimation is an important one, as the estimated gene variances have a key impact on the quality of the results. For the sake of simplicity we have limited ourselves here to a simple procedure, involving sample estimates of gene variances from pre-adjusted data. More sophisticated approaches will have to be investigated.

References

[1] Patrick Cahan, Felicia Rovegno, Denise Mooney, John C. Newman, George St. Laurent III, and Timothy A. McCaffrey. Meta-analysis of microarray results: challenges, opportunities, and recommendations for standardization. *Gene*, 401(1-2):12–18, 2007.

[2] E.M. Conlon, J.J. Song, and A. Liu. Bayesian meta-analysis models for microarray data: a comparative study. *BMC Bioinformatics*, 8:80, 2007.

[3] Markus W. Covert, Eric M. Knight, Jennifer L. Reed, Markus J. Herrgard, and Bernhard O. Palsson. Integrating high-throughput and computational data elucidates bacterial networks. *Nature*, 429:92, 2004.

[4] F. Hong and R. Breitling. A comparison of meta-analysis methods for detecting differentially expressed genes in microarray experiments. *Bioinformatics*, 24:374, 2008.

[5] Adaikalavan Ramasamy, Adrian Mondry, Chris C. Holmes, and Douglas G. Altman. Key issues in conducting a meta-analysis of gene expression microarray datasets. *PLoS Medicine*, 4:1320–1332, 2009.

[6] C.R. Rao. Estimation of variance and covariance components minque theory. *J. Multivar. Anal.*, 1:257, 1971.

[7] M.-C. Roubaud and B. Torrésani. Approche variationnelle pour la fusion de jeux de données d'expression génique. In *Proceedings of the 41-th french Journées de Statistique*, 2009.

[8] M.-C. Roubaud and B. Torrésani. A new approach for merging gene expression datasets. In *Proceedings of IEEE-SSP, the IEEE workshop on Statistical Signal Processing, Nice*, 2011.

[9] Dinesh Singh, Phillip G. Febbo, Kenneth Ross, Donald G. Jackson, Judith Manola, Christine Ladd, Pablo Tamayo, Andrew A. Renshaw, Anthony V. D'Amico, Jerome P. Richie, Eric S. Lander, Massimo Loda, Philip W. Kantoff, Todd R. Golub, and William R. Sellers. Gene expression correlates of clinical prostate cancer behavior. *Cancer Cell*, 1:203, 2002.

[10] Robert O. Stuart, William Wachsman, Charles C. Berry, Jessica Wang-Rodriguez, Linda Wasserman, Igor Klacansky, Dan Masys, Karen Arden, Steven Goodison, Michael McClelland, Yipeng Wang, Anne Sawyers, Iveta Kalcheva, David Tarin, and Dan Mercola. In silico dissection of cell-type-associated patterns of gene expression in prostate cancer. *Proceedings of the National Academy of Sciences of the USA*, 101:615, 2004.

[11] Grace Wahba. *Spline Models for Observational Data*. SIAM, CBMS-NSF Regional Conference Series in Applied Mathematics, 1990.